

Capacity Machines

Andrew Jones

HEC Development Manager, University of Manchester

SOS10 “Capacity Machines” Panel

8 March 2006, Maui

- A vast amount of scientific and technical computing may be characterized as throughput or capacity computing where the dominant parallelism exploited is at the task or job boundary and there is little or no communication required between tasks except at the file I/O name space. Many large conventional systems perform this class of work well including Grid systems, commodity clusters, and the more tightly coupled MPPs such as SNL/Cray Red Storm and LLNL/IBM Purple. If contrasted with “capability” computing, capacity computing and the systems that serve this role are considered second class and yet their importance to the overall computing challenge is enormous. Systems optimized for this can deliver superior performance to cost to capability machines because of reduced investment in system area networks. Ironically, once such class of system, commodity clusters, comprise the majority of the deployed installed based on the Top-500 List which is intended to document supercomputers and hence capability computing. This panel will examine the current state of capacity computing both from a workload perspective and in terms of the state-of-the art in hardware systems and will consider the current challenges and future directions for this important class of scientific processing.
- Principal factors that distinguish capacity systems from other forms of computing systems?
- Role and impact of capacity computing for current and future scientific problems?
- Technical challenges confront the continued growth of capacity computing performance?
- Dominant directions for future generation capacity computing and system types?
- Will the current division between capacity and capability computing be retained over the next decade or will there emerge a different useful distinction in form and function?

“Capacity”?

- Embarrassingly parallel
- “Research group” systems
- (Very) large quantity of 1-32 way jobs

The model?

- Capability computing is often likened to the “Formula 1” element
 - the fastest, the best - when raw \$ plays a 2nd role to raw power
 - a low volume business
 - at the very leading edge, often even creating the leading edge
- Capacity computing must then be our everyday cars
 - cost effectiveness, reliability, personal performance
 - high volume business, highly commoditised
 - uses developed, proven, stable technology

What we really see in this model ...

- Capability computing = “Formula 1”
 - cost effectiveness, reliability, “personal” performance
 - dominated by “proven” technology/vendors
- Capacity computing = our everyday cars
 - when raw \$ plays more important role than lifetime delivered compute-power per \$
 - at the very leading edge, creating new “integrations of components”, take the build/use/manage pain for lower initial \$

What's broken?

- Is the model broken, or our implementation of it?

Broken model

- Formula 1 & everyday cars is not right
- Perhaps Spacetravel & Airtravel?
- Spacetravel = capability
 - The aspiration, leading the technology and other challenges on behalf of science and engineering and even society
 - Custom build, but high engineering, RAS standards, etc
 - Low volume
- Airtravel = capacity
 - Everyday high volume use by all of us
 - Raw \$
 - BUT no way would we all take a custom “lowest build \$, I built this myself aircraft” back to the mainland!

Broken model

- I cannot find a model (comparator in another industry) that matches the reality of the way we implement the capability/capacity pair and of capacity in particular
- So it must be our implementation that is wrong
 - Or my imagination is limited ...

Broken implementation

- Back to the Formula 1 model
- A F1 engine is designed to last the 72 laps ONLY – an engine that's last longer than this has been over-engineered in strength/weight etc and thus could be tweaked to give better performance (speed/power)
- Capability should be at the leading edge – the absolute maximum performance – only the stability required to complete the race/problem, not to be a full high-RAS service as a bank might require
- It's a proving ground for the technologies (& processes) that will be used for the bulk of the real science (which is done on capacity computing)

The real capacity

- And that's where the fault lies in all the models:
- The overwhelming proportion of science is done on capacity type machines (research group clusters and similar)
 - Some very low volume of high profile science is done on capability
 - Just as most of our travelling is done by car or plane
- So the capacity machines/buys should reflect this need
 - Reliable, stable, simple, works on day one
 - Most (lifetime delivered) cost-effective performance for *my* workload
 - Many vendors are ready for this (industry does this in many cases)
 - it is us as research/academic/govt labs that don't follow the model

What should we do?

- Buying a capacity machine
 - HP, SGI, Cray, IBM P5, Bull, NEC, Sun, etc ???
- Buying a capability machine
 - Any of above when stretched to their limit (scale etc)
 - Great big custom engineered “cluster” (capability work comes with effort ...) (!!!)
- In other words:
 - “proper” supercomputers can be *effectively* used in either capacity or capability modes
 - Perhaps “clusters” can/should *only* be used in capability mode (!!!)



Manchester Computing